Commentary on:
Herschel Mission Planning Software

Jeremy D. Frank
NASA Ames Research Center
Jeremy.D.Frank@nasa.gov

Astronomical Observation Planning and Scheduling

Planning and scheduling for astronomical observatories is a significant challenge, both due to the computational complexity of the underlying problem and due to the engineering challenges of building useful tools. Whether observatories are ground-based, aircraft-based or space-based, the fundamental problem is to select what targets to observe, and when. A myriad of constraints must be respected, and preferences (both in terms of science and observing conditions) impose an optimization criteria on schedules.

The Herschel planning and scheduling problem fits this mold. Herschel is a space-based telescope with several instruments, setup and teardown constraints that influence which instrument and which mode is used on a given observing day, a variety of operational constraints limiting spacecraft pointing and other spacecraft activities, and, of course, observations and supporting activities to schedule. The problem is broken down into long-term (roughly annual) and daily timeframes. The problem is also subdivided by instrument mode, that is, one instrument mode is employed per day.

This paper starts with an abstract description of the Herschel mission and its various constraints, followed by a description of the capabilities of the planning and scheduling software. Considerable attention is paid to the capabilities offered to mission operators tool users; schedule visualizations, geometric perspectives on the pointing constraints, and interfaces to target lists from astronomical object catalogs. Attention is also paid to various report generation functions: query interfaces to databases, schedule status, future schedule statistics, and so on. In this sense, the paper is quite informative, as it describes all of the ways in which the scheduling problem and its solution are visualized by the mission operations community. It is notable that both geometric and timeline based visualizations are presented.

The Mission, the Design, or the Technology?

The value of applications papers lies in the synthesis of the specific application (the mission), the design of the solution, and the technology and theory underlying the design. Unfortunately, this paper focuses largely on the design of the solution at the expense of the other two facets.

Many interesting, if not important, details of the application are not completely specified. A tantalizing list of constraints and preferences is provided in the opening paragraph, but many details are missing: how many 'grades' of proposal are there? How does one trade off slewing, helium optimization and proposal completion? Are there absolute time or relative ordering constraints within proposals that must be respected? Are there long durations for targets (e.g. > 12 hours) that must be split? Do proposals mix observations with multiple instruments? Is there choice in how a specific target is observed, and does the choice introduce the option of lower or higher quality schedules?

The paper is also disappointingly light on technical details concerning the scheduling algorithm. For example, how are instrument modes assigned to days during long term scheduling prior to scheduling individual days? What are the different strategies that can be used to schedule observations? Are there strategies that work well or poorly? How long does scheduling take to run, and are sacrifices needed to ensure scheduling is completed in time? How good are the resulting schedules? When a daily plan is created, how is information from the long-range plan conveyed to the short range plan? How does schedule execution status feed back to short-term or long-term scheduling?

These shortcomings are especially disappointing because, as a result, this paper provides no lessons for those building other schedulers for other observatories to learn from. It is also not clear if the scheduling algorithms used in this tool were influenced by other work performed by the astronomy or computer science communities.

Consider, by contrast, the SPIKE scheduler developed by STSCI; numerous publications (including, but not limited to, [1][2][3]) allow both the academic and space operations community a chance to learn from the SPIKE experience. While SPIKE itself need not have formed the basis of the Herschel scheduler, certainly the deep experience offered by SPIKE could inform the development of Herschel's scheduler. By the same token, Herschel's legacy could inform the next generation of astronomical observatory schedulers.

References

[1] JOHNSTON, P. and MILLER, G. 1994. SPIKE: Intelligent Scheduling of Hubble Space Telescope Observations. In Zweben, M. and Fox, M. Intelligent Scheduling. Morgan Kauffmann, San Francisco, CA. USA pp. 391 – 422.

[2] Giuliano, M. Achieving Stable Observing Schedules in an Unstable World. Astronomical Data Analysis Software and Systems VII, ASP Conference Series, Vol. 145, 1998, Editors: R. Albrecht, R.N. Hook and H. A. Bushouse, pp. 271-274.

[3] Giuliano, M., Rager, R., and Ferdous, N. Towards a Heuristic for Scheduling the James Webb Space Telescope. Proceedings of the 17^{th} International Conference on Automated Planning and Scheduling, 2007.